



THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant : Massimiliano Antonio Poletto et al. Art Unit : 2184
Serial No. : 10/066,232 Examiner : Perungavoor, Venkatanaray
Filed : January 31, 2002
Title : DENIAL OF SERVICE ATTACKS CHARACTERIZATION

Mail Stop Appeal Brief - Patents

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF ON BEHALF OF MASSIMILIANO ANTONIO POLETTO ET AL.

The Appeal Brief fee has already been paid. Please apply any other charges or credits to
Deposit Account No. 06-1050.

CERTIFICATE OF MAILING BY FIRST CLASS MAIL

I hereby certify under 37 CFR §1.8(a) that this correspondence is being deposited with the United States Postal Service as first class mail with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Date of Deposit

November 2, 2006

Signature

Maie Collins

Typed or Printed Name of Person Signing Certificate

Maie Collins

(i.) Real Party In Interest

The real party in interest in the above application is Mazu Networks, Inc.

(ii.) Related Appeals and Interferences

The appellant is not aware of any appeals or interferences related to the above-identified patent application.

(iii.) Status of Claims

This is an appeal from the decision of the Primary Examiner in an Office Action dated August 4, 2006 rejecting claims 1-40. The claims have been twice rejected. Claims 1-40 are the subject of this appeal.

(iv.) Status of Amendments

All amendments have been entered. Appellant filed a Request for Reconsideration and a Notice of Appeal on February 7, 2006. Appellant subsequently filed an Appeal Brief on May 26, 2006. In response, the examiner furnished the above mentioned office action, from which Appellant now appeals.

(v.) Summary of Claimed Subject Matter

Background

The invention relates to techniques to thwart network-related denial of service attacks. In denial of service attacks, an attacker sends a large volume of malicious traffic to a victim in an attempt to prevent the victim from responding to legitimate traffic.

Appellant's Invention

Claim 1

One aspect of Appellant's invention is set out in claim 1, as a process that monitors network traffic through a monitoring device "The gateway 26 devices are located at the edges of the Internet 14, for instance, at the entry points of data centers. The gateway devices constantly

analyze traffic, looking for congestion or traffic levels that indicate the onset of a DoS attack.” [Specification page 5, lines 23-27] disposed between a data center and a network for thwarting denial of service attacks on the data center. “The victim 12 is coupled to the Internet 14 or other network. For example, the victim 12 has a web server located at a data center (not shown).” [Specification page 4, lines 30-32].

Inventive features of claim 1 include a detection process to determine if the values of a parameter of network traffic exceed normal values for the parameter to indicate an attack on the data center. “Several methods can be used separately or in combination to detect, malicious traffic flows. For example, the gateway 26 can detect DoS attacks using at least one or more of the following methods including:” [Specification page 17, lines 23-26].

Inventive features of claim 1 also a characterization process to build a histogram for the parameter to compute significant outliers in a parameter and classify the attack. “Referring to FIG. 12, attack characterization 139 is based on comparison of historical histogram data with near-real-time histogram data for one or several parameters (e.g., source/dest IP address, source/dest TCP/UDP ports, IP protocol, IP TTL, IP length, hash of payload fragment; IP TOS field and TCP flags). [Specification page 25, lines 23-29].

Typically, historical histograms are based on time periods that can range from 1 hour to 1 week. During an attack, attack histograms are produced 142 for time periods, e.g., in the 10-300sec range. For each parameter, the two histograms are normalized 144 (integral set equal to 1) and their difference 146 is used to compute 148 significant outliers.” [Specification page 25, line 29 to page 26, line 3].

Inventive features of claim 1 also include a filtering process for filtering of network packets based on the characterization process. “The attack characterization process 139 correlates 150 the suspicious parameters and determines existence of correlations of those parameters that can be indications of attacks. If under attack 152 the process 139 will employ filtering.” [Specification page 26, lines 19-23].

Claim 7

Claim 7 claims a method for thwarting denial of service attacks on a data center. This feature generally finds support at least as the analogous feature of claim 1.

Inventive features of claim 7 include producing a histogram of received network traffic for at least one parameter of network packets. This feature generally finds support at least as the analogous feature of claim 1.

Inventive features of claim 7 also include characterizing an attack based on comparison of a historical histogram with the produced histogram data for one or more parameters. This feature generally finds support at least as the analogous feature of claim 1.

Claim 21

Claim 21 claims a monitoring device for thwarting denial of service attacks on a data center. "Another alternate implementation could combine thresholds with a histogram analysis, and trigger traffic characterization whenever a histogram for some parameter differed significantly (by uniformity test, or for example, by subtracting normalized histograms) from the historical histogram. [Specification page 19, lines 10-15]. ... Optionally, the gateway 26 executing the detection process 131 can build 134 a histogram ... for any attribute or function of an attribute of network packets to use in determining if an attack is occurring." [Specification page 25, lines 12-16].

Inventive features of claim 21 include a computing device executing a process to build at least one histogram for at least one parameter of network traffic. This feature generally finds support at least as the analogous feature of claim 1.

Inventive features of claim 21 also include a process to characterize an attack based on a comparison of a historical histogram of the at least one parameter to the built at least one histogram for the at least one parameter. This feature generally finds support at least as the analogous feature of claim 1.

Claim 28

Claim 28 claims a computer program product residing on a computer readable medium. "The gateway 26 comprises a software program executing on a device, e.g., a computer 27 that is disposed at the edge of the data center 20 behind an edge router coupled in the Internet 14." [Specification page 6, lines 26-28].

Inventive features of claim 28 include instructions to build a histogram for a parameter of network traffic. This feature generally finds support at least as the analogous feature of claim 1.

Inventive features of claim 28 also include instructions to use the histogram data for the parameter to characterize an attack. This feature generally finds support at least as the analogous feature of claim 1.

Claim 32

Claim 32 claims a method of protecting a data center during a denial of service attack. This feature generally finds support at least as the analogous feature of claim 1.

Inventive features of claim 32 include monitoring network traffic through a gateway disposed between the data center and a network. "During normal operation, the gateway collects 132 information about normal network traffic for these parameters. The gateway determines 137 normal or reasonable values for each of the attributes or functions of attributes that the gateway tracks, as mentioned above. The detection process 131 uses some statistic about the flow, such as an average of a traffic ratio of the standard deviation to a mean of a histogram attribute over a time window, e.g., over a minute, or hour, etc." [Specification page 24, lines 10-18.

Inventive features of claim 32 also include determining if values of at least one parameter exceed normal, threshold values expected for the parameter to indicate an attack on the site. "The gateway determines 137 normal or reasonable values for each of the attributes or functions of attributes that the gateway tracks, as mentioned above." [Specification page 24, lines 12-14].

Inventive features of claim 32 also include producing a histogram for the at least one parameter of network traffic to characterize the attack by comparing the histogram to at least one historical histogram for that parameter. "Consider a historical histogram H_i , and a current (under attack) histogram C_i . The use of the noise reduction process 151 is to normalize each histogram so the integral of each histogram equals one. Then for each bucket i component of the histogram, the noise reduction process computes a difference value D_i ($D_i = C_i - H_i$) to determine a difference value for each bucket relative to historical norm, and produces a "difference histogram," D as generally described above to find outliers." [Specification page 27, line 5-13].

Inventive features of claim 32 also include filtering out traffic based on characterizing the traffic, which the gateway deems to be part of an attack. This feature generally finds support at least as the analogous feature of claim 1.

Claim 37

Claim 37 is directed to a method to reduce blocking of legitimate traffic in a process to protect a victim site during a denial of service attack. This feature generally finds support at least as the analogous feature of claim 1.

Inventive features of claim 37 include producing a histogram of network traffic to characterize an attack. This feature generally finds support at least as the analogous feature of claim 1. This feature generally finds support at least as the analogous feature of claim 1.

Inventive features of claim 37 also include filtering out traffic deemed part of an attack. This feature generally finds support at least as the analogous feature of claim 1. "Referring to FIG. 14, a second filtering mechanism is aggregate filtering 170. Aggregate filtering 170 allows constant-time filtering, independent of the number of individual parameter values used for filtering. Based on the correlation histogram, the process 170 constructs 172 a master correlation vector (a bit vector that has 1-bits corresponding to the most important parameter correlations)." [Specification page 29, lines 26-32].

Inventive features of claim 37 also include constructing a master correlation vector having asserted bits corresponding to the most important parameter correlations. "Based on the correlation histogram, the process 170 constructs 172 a master correlation vector (a bit vector that has 1-bits corresponding to the most important parameter correlations)." [Specification page 29, lines 29-31].

Inventive features of claim 37 also include initializing a packet's correlation bit vector to 0, and for every parameter: "Given a packet, the process initializes 174 the packet's correlation bit vector to 0. The process 170 loops for every parameter (TTL, etc.), and retrieves 176 the parameter in the parameter suspicious vector to construct 178 the packet's correlation bit vector. If the bit in the suspicious vector is 1, the process sets the relevant bit in the packet's correlation vector to a 1 (in the example, bit 0 for TTL, 1 for source address, 2 for dest address)." [Specification page 30, lines 4-12].

Inventive features of claim 37 also include retrieving the parameter in a parameter suspicious vector to construct the packet' correlation bit vector. "In the example above, since buckets 4 and 5 were deemed suspicious, the master correlation bit vector would be 00110000 (bits 4 and 5, decimal 48)." [Specification page 30, lines 1-3].

Inventive features of claim 37 also include using the value of the packet's correlation bit vector to index into the master correlation bit vector. "The process uses 180 the value of the packet's correlation bit vector to index into the master correlation bit vector. The process 170 tests 182 the indexed bit in the master correlation vector. If the bit in the master correlation bit vector is a one, the packet is dropped, otherwise the packet is forwarded." [Specification page 30, line 13-18].

(vi.) Grounds of Rejection to be Reviewed on Appeal

1. Claims 7, 9-14, 19-23, 26 stand rejected under 35 U.S.C. 102(e) as being anticipated by U.S. Patent 6,769, 066 B1 to Botros et al. (Botros).

2. Claims 1-6, 8, 15-18, 24-25, 27, 30-40 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Botros et al. in view of U.S. Patent Application 2002/0107960 A1 to Wetherall et al.(Wetherall).

(vii.) Argument

Anticipation

"It is well settled that anticipation under 35 U.S.C. §102 requires the presence in a single reference of all of the elements of a claimed invention." *Ex parte Chopra*, 229 U.S.P.Q. 230, 231 (BPA&I 1985) and cases cited.

"Anticipation requires the presence in a single prior art disclosure of all elements of a claimed invention arranged as in the claim." *Connell v. Sears, Roebuck & Co.*, 220 U.S.P.Q. 193, 198 (Fed. Cir. 1983).

"This court has repeatedly stated that the defense of lack of novelty (i.e., 'anticipation') can only be established by a single prior art reference which discloses each and every element of

the claimed invention." *Structural Rubber Prod. Co. v. Park Rubber Co.*, 223 U.S.P.Q. 1264, 1270 (Fed. Cir. 1984), citing five prior Federal Circuit decisions since 1983 including *Connell*.

In a later analogous case the Court of Appeals for the Federal Circuit again applied this rule in reversing a denial of a motion for judgment n.o.v. after a jury finding that claims were anticipated. *Jamesbury Corp. v. Litton Industrial Prod., Inc.*, 225 U.S.P.Q. 253 (Fed. Cir. 1985).

After quoting from *Connell*, "Anticipation requires the presence in a single prior art disclosure of all elements of a claimed invention arranged as in the claim," 225 U.S.P.Q. at 256, the court observed that the patentee accomplished a constant tight contact in a ball valve by a lip on the seal or ring which interferes with the placement of the ball. The lip protruded into the area where the ball will be placed and was thus deflected after the ball was assembled into the valve. Because of this constant pressure, the patented valve was described as providing a particularly good seal when regulating a low pressure stream. The court quoted with approval from a 1967 Court of Claims decision adopting the opinion of then Commissioner and later Judge Donald E. Lane:

[T]he term "engaging the ball" recited in claims 7 and 8 means that the lip contacts the ball with sufficient force to provide a fluid tight seal ***** The Saunders flange or lip only sealingly engages the ball 1 on the upstream side when the fluid pressure forces the lip against the ball and never sealingly engages the ball on the downstream side because there is no fluid pressure there to force the lip against the ball. The Saunders sealing ring provides a compression type of seal which depends upon the ball pressing into the material of the ring. *** The seal of Saunders depends primarily on the contact between the ball and the body of the sealing ring, and the flange or lip sealingly contacts the ball on the upstream side when the fluid pressure increases. 225 U.S.P.Q. at 258.

Relying on *Jamesbury*, the ITC said, "Anticipation requires looking at a reference, and comparing the disclosure of the reference with the claims of the patent in suit. A claimed device is anticipated if a single prior art reference discloses all the elements of the claimed invention as arranged in the claim." *In re Certain Floppy Disk Drives and Components Thereof*, 227 U.S.P.Q. 982, 985 (U.S. ITC 1985).

Obviousness

"It is well established that the burden is on the PTO to establish a prima facie showing of obviousness, *In re Fritsch*, 972 F.2d 1260, 23 U.S.P.Q.2d 1780 (C.C.P.A., 1972)."

"It is well established that there must be some logical reason apparent from the evidence or record to justify combination or modification of references. *In re Regal*, 526 F.2d 1399 188, U.S.P.Q.2d 136 (C.C.P.A. 1975). In addition, even if all of the elements of claims are disclosed in various prior art references, the claimed invention taken as a whole cannot be said to be obvious without some reason given in the prior art why one of ordinary skill in the art would have been prompted to combine the teachings of the references to arrive at the claimed invention. *Id.* Even if the cited references show the various elements suggested by the Examiner in order to support a conclusion that it would have been obvious to combine the cited references, the references must either expressly or impliedly suggest the claimed combination or the Examiner must present a convincing line of reasoning as to why one skilled in the art would have found the claimed invention obvious in light of the teachings of the references. *Ex Parte Clapp*, 227 U.S.P.Q.2d 972, 973 (Board. Pat. App. & Inf. 985)."

"The mere fact that the prior art could be so modified would not have made the modification obvious unless the prior art suggested the desirability of the modification." *In re Gordon*, 221 U.S.P.Q. 1125, 1127 (Fed. Cir. 1984).

Although the Commissioner suggests that [the structure in the primary prior art reference] could readily be modified to form the [claimed] structure, "[t]he mere fact that the prior art could be so modified would not have made the modification obvious unless the prior art suggested the desirability of the modification." *In re Laskowski*, 10 U.S.P.Q. 2d 1397, 1398 (Fed. Cir. 1989).

"The claimed invention must be considered as a whole, and the question is whether there is something in the prior art as a whole to suggest the desirability, and thus the obviousness, of making the combination." *Lindemann Maschinenfabrik GMBH v. American Hoist & Derrick*, 221 U.S.P.Q. 481, 488 (Fed. Cir. 1984).

Obviousness cannot be established by combining the teachings of the prior art to produce the claimed invention, absent some teaching or suggestion supporting the combination. Under Section 103, teachings of references can be combined only if there is some suggestion or incentive to do so. *ACS Hospital Systems, Inc. v. Montefiore Hospital*, 221 U.S.P.Q. 929, 933 (Fed. Cir. 1984) (emphasis in original, footnotes omitted).

"The critical inquiry is whether 'there is something in the prior art as a whole to suggest the desirability, and thus the obviousness, of making the combination.'" *Fromson v. Advance Offset Plate, Inc.*, 225 U.S.P.Q. 26, 31 (Fed. Cir. 1985).

1. Claims 7, 9-14, 19-23, 26 are patentable over Botros.

Claims 7, 19, 20 and 21 and 26

For the purposes of this appeal only, Claims 7 and 21 stand or fall together. Claim 7 is representative of this group of claims.

Claim 7 is directed to a method for thwarting denial of service attacks on a data center. Claim 7 is neither anticipated nor obvious over Botros, since Botros neither describes nor suggests the features of producing a histogram of received network traffic for at least one parameter of network packets and characterizing an attack based on comparison of a historical histogram with the produced histogram data for one or more parameters.

The examiner contends that: "... Botros discloses the producing of histogram of received network traffic based on parameters see Col 10 Ln 40-53 & Fig. 14 item 1402; characterization of attack based on comparison of historical histogram of data for parameter see Col 8 Ln 28-39 & Fig. 14 item 1404."

Appellant contends that Botros fails to disclose at Col 10 Ln 40-53 & Fig. 14 item 1402 or and Col 8 Ln 28-39 & Fig. 14 item 1404 or elsewhere the features of claim 7. At Col. 10, lines 40-53, Botros discloses:

In step 904 a histogram or density graph is defined for the normal data gathered at step 902. An example of a histogram for normal data is shown in FIG. 10. In the described embodiment the feature values are normalized to a value

between -5 and +5. Generally, most normal behavior for an activity will have a normalized feature value close to the zero value range, indicating normal or non-anomalous behavior. Anomalous behavior for a particular feature has values closer to -5 or +5 depending on the activity. Generally, a normalized feature value closer to -5 indicates that the particular activity is being performed less frequently than normal and a value closer to +5 indicates the opposite. Characteristics of the histogram are described in greater detail in FIG. 10.

While Botros discloses a histogram, Botros fails to disclose producing a histogram of received network traffic. In the passage relied on by the examiner, Botros describes normalizing data in a histogram. However, neither in that passage nor elsewhere does Botros disclose that the histogram is of received network traffic.

To the contrary, Botros discloses that the histogram is derived from "user logs." While, Botros discloses raw user data, that data is used to produce user logs. Botros discloses that: "These files contain raw user data generated from various system resources and, in the described embodiment, are parsed and organized according to user and time of activity." [Botros, Col.5, lines 47-50]. User data is further defined by Botros, as:

FIG. 3 is a schematic diagram showing the formation of user activity files 12, or the raw user data, in accordance with one embodiment of the present invention. As mentioned above, user activity files 12 contain raw data of activities performed by users. As described below, user activity files 12 are made up of numerous individual user logs, such as user log 204 of FIG. 3. In the described embodiment, the users are on one particular computer system, typically supported by a mainframe computer and operating system. In other embodiments, the raw data can come from several computer systems each supported by different computers. Similarly, score 110 can be derived from data from one or more computer systems and can measure potential intrusions for one or all systems. A computer system 200 is shown containing a number of sources from which raw user activity data is drawn. Examples of these sources or files include operating system files containing executed commands, operations on programs, exceptions, operations on files, and other more data-specific files such as badge-in data. In the described embodiment the sources are maintained by the Multiple Virtual Storage ("MVS") operating system of the IBM Corporation, and used on IBM mainframe computers. These data sources are part of the MVS operating system and are created and maintained as part of the operating system. The process can be used in computer systems using operating systems other than MVS such as a Unix-based operating system. Using the example from above, to determine the time between login failures, the intrusion program checks user activity files 12. [Botros, Col.6, lines 24-52]

Clearly, Botros does not disclose producing a histogram based on received network traffic. Botros produces a histogram based on user or user's activity and specifically activity on a computer system.

Moreover, Appellant contends that were Botros's teachings be construed to meet the claimed language of producing a histogram of received network traffic, Botros still fails to describe: "characterizing an attack based on comparison of a historical histogram with the produced histogram data for one or more parameters." At Col 8 Ln 28-39 & Fig. 14 item 1404, Botros discloses

The peer historical standard deviation can be used to assign various weightings to the peer historical mean based on several criteria, such as time or other factors in the system. For example, a peer historical mean calculated four months prior to the present can be assigned a lighter weight than the historical mean calculated two days prior to the present with regard to determining the standard deviation. This is based on the assumption that behavior from two days ago should be given more importance than behavior from four months ago. In another example, a higher or lower weight can be assigned based on particular days of the weeks.

However, this teaching does not describe what the examiner uses it for, namely to anticipate the feature of: "characterizing an attack based on comparison of a historical histogram with the produced histogram data for one or more parameters." Rather, this teaching is directed to assigning weightings to factors in the computer system, not to characterization of an attack based on a comparison of historical and produced histograms. Moreover, when Botros is view as a whole, it is clear that this teaching is not related at all to Appellant's claimed characterizing but instead is directed to training a model. For example, immediately following that excerpt, Botros describes: "At step 512 the intrusion detection program determines whether there are any other activities from the predetermined list of activities to be examined. If so, control returns to step 504 where another activity is selected and the process is repeated. If there are no more activities, the process of generating peer historical data is complete." Thus, these teachings are merely directed at generating peer historical data, which is used by Botros in the Neural Network Training discussed in Col. 10.

Nonetheless Botros neither describes nor suggests: "characterizing an attack based on comparison of a historical histogram with the produced histogram data for one or more parameters.", because Botros does not use the histograms to characterize an attack but instead uses a neural network that is trained by histograms.

Moreover, Botros does not teach to use a comparison of an historical histogram to the produced histogram. Rather, as Botros clearly describes:

User activity files 12 and historical data 102 are used as input to a feature generator or builder 104. In the described embodiment, feature generator 104 is implemented involving an equation for calculating a time-weighted mean, discussed in greater detail in FIGS. 6 and 7. The output from feature generator 104 is a features list 106. In the described embodiment, features list 106 contains 47 features which can be classified into several different categories such as violations, user activities, computer and network loads, and so on. Characteristics of feature list 106 are described in greater detail in FIG. 8. Individual features from features list 106 are used as input to a model 108. As is well known in the field of computer science, there are many different model processes, such as linear regression, Markov models, graphical models, and regression models. A model is trained to evaluate features to recognize the possibility of a network intrusion. By training model 108 to process certain types of features, it can recognize potential intrusions. As is well known in the art, a model can accept different types of features. One example of a feature is user login failure, such as the time between login failures for a particular user. Once the model receives all input features, it calculates a score 110. This score is based upon the input features and how the model has been trained. In the described embodiment, the model is trained using a neural network algorithm. A score 110 can be normalized to a number between 0 and 1000, a high number indicating a stronger possibility of an intrusion. An advantageous method for training a suitable model is discussed in FIGS. 9 through 14." [Botros, Col.5, line 62 to col. 6, line 24].

Thus, while Appellant characterizes an attack based on comparison of a historical histogram with the produced histogram data for one or more parameters, Botros uses a model that relies on "features," e.g., "In the described embodiment, features list 106 contains 47 features which can be classified into several different categories such as violations, user activities, computer and network loads, and so on. Characteristics of feature list 106 are described in greater detail in FIG. 8. Individual features from features list 106 are (sic) used as input to a model 108." [Botros col. 5, line 63 to col. 6, line 6].

Appellant's characterizing based on comparison of a historical histogram with the produced histogram data for one or more parameters thus is not described by Botros model features list.

Therefore Claim 7 and claim 21 are allowable over Botros since "Anticipation requires the presence in a single prior art disclosure of all elements of a claimed invention arranged as in the claim." *Connell v. Sears, Roebuck & Co.*, 220 U.S.P.Q. 193, 198 (Fed. Cir. 1983) and Botros fails to disclose the features of producing a histogram of received network traffic for at least one

parameter of network packets and characterizing an attack based on comparison of a historical histogram with the produced histogram data for one or more parameters.

Claims 9 and 10

For the purpose of this appeal only, claims 9 and 10 each distinguish over the reference.

Claim 9 recites that the historical histogram is based on time periods that can range from 1 hour to 1 week or more, whereas, claim 10 recites that the produced histogram is produced during an attack and over time periods of about 10-300 sec or so.

The examiner contends that: "Regarding Claim 9 and 10, Botros discloses the varying of time to get an accurate historical data see Col 7 Ln 24-57.

At the outset Appellant has not claimed "varying of time to get an accurate historical data." Moreover, Col. 7 lines 24-57 does not deal with the claimed features time periods of the historical or the produced histograms but rather a count of the number of time an event occurs. Accordingly each of claims 9 and 10 distinguish over Botros.

Claims 11-14

For the purpose of this appeal only, claims 11-14 stand or fall together. Claim 11 is representative of this set of claims.

The examiner contends that: "Regarding Claim 11, 23, Botros discloses the normalizing of the histograms and computing the difference to significant outlier of suspicious traffic see Col 9 Ln 24-50."

Claim 11, limits claim 7, and includes normalizing the produced and the historical histograms for each parameter and computing their difference to identify significant outliers that are considered indicators of suspicious traffic.

At no point does Botros disclose any computation of a difference in historical and produced histograms. Botros as discussed above uses histograms to train a model not to computer significant outliers.

Claims 19, 20 and 26

Claim 19 calls for the method being executed on a data collector and claim 20 calls for the method being executed on a gateway.

The examiner states: "Regarding Claim 19 and 20, 26, Botros discloses the data collector see Fig. 3 item 202 and gateway see item 200."

Neither 200 nor 202 are the claimed data collector or gateway. Item 202 is a log, e.g., a database of some sort and item 200 is a computing device. Neither the log nor the computer is disclosed as being the claimed data collector or gateway. Indeed, the log is not even disclosed as executing anything.

Claims 22 and 23

Regarding Claim 22, Claim 22 limits the monitoring device of claim 21. Appellant contends that Botros fails to disclose the monitoring device including a process to correlate suspicious parameters to reduce blocking of legitimate traffic. The examiner contends that: "Regarding Claim 12 and 14, 22, Botros discloses the correlation process that correlates the parameters and indicates the types of attacks see Col 13 Ln 24-41.

Neither at that passage nor elsewhere does Botros correlate suspicious parameters to reduce blocking of legitimate traffic. Rather, Botros only teaches to produce a score for the anomalous or normal features list in the context of training the genetic model. Nowhere does Botros suggest much less describe: "a process to correlate suspicious parameters to reduce blocking of legitimate traffic."

**2. Claims 1-6, 8, 15-18, 24-25, 27, 30-40
patentable over Botros et al. in view of
Wetherall.**

Claims 1, 3 and 4

For the purposes of this appeal only Claims 1, 3 and 4 stand or fall together. Claim 1 is representative of this group of claims.

Claim 1 recites a process that monitors network traffic through a monitoring device disposed between a data center and a network for thwarting denial of service attacks on the data

center. Claim 1 includes the features of a detection process to determine if the values of a parameter of network traffic exceed normal values for the parameter to indicate an attack on the data center and a characterization process to build a histogram for the parameter to compute significant outliers in a parameter and classify the attack. Claim 1 also includes a filtering process for filtering of network packets based on the characterization process.

The examiner contends that:

13. Regarding Claim 1, 28, Botros discloses the monitoring of device for attacks where detecting process to determine if the parameter of network traffic exceed normal value see Col 10 Ln 40-53 & Col 9 Ln 1-23; a process to build histogram for the parameter see Fig. 9 item 904 & Fig. 10. But does not disclose a filtering of network packets based on the characterization process using histogram for parameter to compute significant outliers in a parameter and classify the attack. However, Wetherall discloses the filtering processes based on the characterization process using histogram for parameter to compute significant outliers in a parameter and classify the attack see Fig. 2 item 206, 208, 210, 212. It would be obvious to one having ordinary skill in the art at the time of the invention to include the filtering processes based on the characterization process using histogram in the invention of Botros in order to route out the undesirable packets as taught in Wetherall see Par. 0041.

Appellant contends that no combination of Botros with Wetherall would suggest the features of claim 1. Contrary to the examiner's contentions, Botros neither describes nor suggests at least: "a process to determine if the parameter of network traffic exceed normal value," whether at Col 10 Ln 40-53 or Col 9 Ln 1-23. Botros also does not suggest "a process to build histogram for the parameter see Fig. 9 item 904 & Fig. 10." Rather as discussed above, Botros discloses the use of histograms of user data to train a model. The user data are not values of a parameter of network traffic, nor are the data used to indicate if the values exceed normal values for the parameter to indicate an attack Botros, also as discussed above, fails to disclose a characterization process to build a histogram for the parameter to compute significant outliers in a parameter and classify the attack

Botros fails to suggest a process that monitors network traffic through a monitoring device disposed between a data center and a network for thwarting denial of service attacks on the data center and a detection process to determine if the values of a parameter exceed normal values for the parameter to indicate an attack on the data center.

The examiner acknowledges that Botros does not disclose the claimed filtering but contends that Wetherall teaches: "... the filtering processes based on the characterization process using histogram for parameter to compute significant outliers in a parameter and classify the attack see Fig. 2 item 206, 208, 210, 212." Appellant disagrees. While Appellant notes that Wetherall does disclose filtering based on a distribution profile of source address, Wetherall fails to disclose any of the specific claimed features of filtering of network packets based on the characterization process.

Claim 1 specifically recites: "a characterization process to build a histogram for the parameter to compute significant outliers in a parameter and classify the attack; and a filtering process for filtering of network packets based on the characterization process." Wetherall does not specifically build a histogram for the parameter to compute significant outliers and thus classify the attack. Rather, Wetherall at [0035] discloses: "At block 204, the gathered and cached source address instances of the packets routed are reported, e.g. to director 102." Wetherall thereafter discusses that: "In any event, at block 206, a spatial, a destination source address range, a migration, and/or a timing profile is constructed (e.g. by director 102) for each of the reported source addresses. At block 208, a determination is made (e.g. by director 102), based at least in part on the constructed (S/D/M/T) profile, on whether any of the reported source addresses should be deemed as having spoof source address instances." Thus, Wetherall teaches to report cached source address, construct a profile and determine whether any of the reported source addresses are spoofed. Wetherall says nothing about using the histogram in a characterization to compute significant outliers in the source address and thus classify the attack. Therefore Wetherall also does not disclose a filtering process for filtering of network packets based on the characterization process.

Appellant further contends that there is no suggestion or motivation to combine Botros with Wetherall. The examiner contends that: "It would be obvious to one having ordinary skill in the art at the time of the invention to include the filtering processes based on the characterization process using histogram in the invention of Botros in order to route out the undesirable packets as taught in Wetherall see Par. 0041." Botros however deals with files that "contain raw user data generated from various system resources" and are described as

“individual user logs... on one particular computer system... . In other embodiments ... from several computer systems each supported by different computers.” [See generally Col. 6, lines 24-54].

Appellant disagrees. Modifying Botros with Wetherall would not serve any purpose in Botros, since Botros does not specifically deal with the problem addressed by Wetherall and filtering spoofed source addresses, as argued by the examiner would not have any effect on the user logs disclosed by Botros at least because the user logs of Botros are not disclosed as being dependent on source address. Accordingly, no combination of these references suggests the features of claim 1.

Claims 2 and 5

For the purposes of this appeal only Claims 2 and 5 stand or fall together. Claim 2 is representative of this group of claims.

Claim 2 limits the characterization process of claim 1 to include the feature that suspicious parameter values are represented by a bit vector with a 1 in every position corresponding to a "bad" value, and a 0 in every position corresponding to a "good" value.

The examiner contends that: “Regarding Claim 2, Botros discloses the vector having bad and good values see Col 9 Ln 51-Col 10 Ln 3.”

Clearly, if Botros does not suggest the claimed filtering process of base claim 1, Botros would not suggest a bit vector or the feature that suspicious parameter values are represented by a bit vector, whether at Col 9 Ln 51-Col 10 Ln 3 or indeed elsewhere. Botros discusses a list of scores, e.g., features values. The features values however are not a bit vector. Neither the teachings at Col 9 Ln 51-Col 10 Ln 3 nor Figure 8 disclosed the bit vector features of claim 2.

Similarly in Wetherall, no such teachings exist with respect to how filtering is specifically accomplished with the bit vectors, as claimed.

Claim 6

Claim 6 further limits the process of claim 1 to parameters that include at least one of source IP address, destination IP address, source TCP/UDP ports, destination TCP/UDP ports, IP protocol, IP TTL, IP length, hash of payload fragment, IP TOS field, and TCP flags.

The examiner contends that: "17. Regarding Claim 5 and 6, 29, Botros does not disclose the aggregate filtering. However, Wetherall discloses the aggregate filtering see Par. 0041 (blanket filtering) and source IP address see Fig. 2 item 208 & 212. For motivation to combine see above Claim 1.

Appellant disagrees. Appellant's claim 6 does require that the parameters include at least one of source IP address However, claim 6 uses that at least one parameter to determine if the values of that parameter of network traffic exceed normal values for the parameter to indicate an attack on the data center and incorporates a characterization process that builds a histogram for the parameter to compute significant outliers in the parameter to classify the attack. Wetherall, deals with Spoofing Attacks, but does not classify the attack based on the source IP Address, but instead at that passage uses a histogram as a basis for filtering of the packets. For instance Wetherall discloses:

Skipping briefly to FIG. 13a-13d and FIG. 14a-14d. FIG. 13a-13b illustrate one each of an example spatial and an example "destination" distribution profile of a source address having spoof instances. Experience has shown that if spoof source addresses are employed in a denial of service attacks against a network node, it is likely that the source addresses will be simultaneously observed in multiple domains of network 100, even domains that are geographically dispersed, as illustrated by the histogram of FIG. 13a.

Claim 8

Claim 8 further limits claim 7 and recites: "filtering network packets sent to the data center based on whether or not a value of the attribute represented in the current histogram is within a normal range of values for the attribute, as determined by comparison to the historical histogram."

The examiner contends that: "Regarding Claim 8, 30-31, Botros discloses the comparison of historical data for ranges see Col 10 Ln 40-53, but does not disclose the filtering process. However, Wetherall discloses the filtering processes based on the characterization process using histogram see Fig. 2 item 206, 208, 210, 212. It would be obvious to one having ordinary skill in the art at the time of the invention to include the filtering processes based on the characterization process using histogram in the invention of Botros in order to route out the undesirable packets as taught in Wetherall see Par. 0041."

As was argued above for claim 2, Botros does not suggest a bit vector or that suspicious parameter values are represented by a bit vector, at Col 9 Ln 51-Col 10 Ln 3 or indeed elsewhere, at least because the list of scores, e.g., features values, are not a bit vector.

The examiner now takes the position that "... Botros discloses the comparison of historical data for ranges see Col 10 Ln 40-53, but does not disclose the filtering process." However, as already argued, Botros does not suggest, much less disclose the claimed bit vector. The examiner argues that "... Wetherall discloses the filtering processes based on the characterization process using histogram see Fig. 2 item 206, 208, 210, 212.", but does not address all of the features of this claim, namely the bit vector. Accordingly, assuming that it is suggested to combine the references, no combination of these references suggests all of the features of claim 8.

Claims 15-17

For the purposes of this appeal only, Claims 15-17 stand or fall together. Claim 15 is representative of this group of claims.

Appellant's claim 15, limits claim 11 and specifically filtering based on attribute, by producing a master correlation vector from a stream of sampled packets and examining the network packets using a process that is constant-time, independently of the number of correlations or of the number of suspicious values for a parameter.

Neither Botros nor Wetherall suggest a bit vector and therefore no combination of these references would suggest the features of producing a master correlation vector from a stream of sampled packets and examining the network packets using a process that is constant-time, independently of the number of correlations or of the number of suspicious values for a parameter. The examiner contends:

Regarding Claim 15-18, 24, 27, 38-40, Botros does not disclose the producing of a vector that is constant and constructing a vector for packets to test whether to forward for drop packets from source address. However, Wetherall discloses the producing of a vector that is constant and constructing a vector for packets test whether to forward for drop see Par. 0056. It would be obvious to one having ordinary skill in the art at the time of the invention to include the filtering processes based on the characterization process using histogram in the invention of Botros in order to route out the undesirable packets as taught in Wetherall see Par. 0041.

Appellant contends that Wetherall does not disclose the features of constructing a vector and does not disclose a master correlation bit vector.

Claim 18

Claim 18, which recites that the attributes include at least one of source IP address, destination IP address, source TCP/UDP ports, destination TCP/UDP ports, IP protocol, IP TTL, IP length, hash of payload fragment, IP TOS field, and TCP flags, is allowable for the reasons discussed in claim 6.

Claim 25

Claim 25 recites dynamically installing of filters on nearby routers, which as the examiner admits is not disclosed by Botros. The examiner uses Wetherall to teach this feature. Specifically the examiner contends that Wetherall: “discloses the installing of filters on routers see Fig. 1 item 106d & Fig. 2 item 210.” However, as argued above there is no suggestion to combine Botros with Wetherall, since Botros does not deal with management of the network, but instead is directed to management of resources on a user's computer.

Claim 32, 35 and 36

For the purposes of this appeal only, Claims 32, 35 and 36 stand or fall together. Claim 32 is representative of this group of claims.

Claim 32 is directed to a method of protecting a data center during a denial of service attack. Features of claim 32 include monitoring network traffic through a gateway disposed between the data center and a network, determining if values of at least one parameter exceed normal, threshold values expected for the parameter to indicate an attack on the site, producing a histogram for the at least one parameter of network traffic to characterize the attack by comparing the histogram to at least one historical histogram for that parameter, and filtering out traffic based on characterizing the traffic, which the gateway deems to be part of an attack.

The examiner contends that: “Regarding Claim 32, Botros discloses the monitoring of device for attacks where detecting process to determine if the parameter of network traffic exceed normal value see Col 10 Ln 40-53 & Col 9 Ln 1-23 through a gateway see Fig. 2 item

104; a process to build histogram for the parameter see Fig. 9 item 904 & Fig. 10 and comparing it with a historical histogram for a parameter see Col 8 Ln 28-39 & Fig. 14 item 1404."

Appellant refers the Board to the above discussion where it is plainly evident that Botros does not disclose any of these features.

The examiner admits that "[Botros] ... does not disclose a filtering of network packets based on the characterization process using histogram." The examiner relies on Wetherall to teach "... the filtering processes based on the characterization process using histogram see Fig. 2 item 206, 208, 210, 21." While Wetherall does disclose filtering, Wetherall does not do so based on a characterization process using a histogram. Wetherall discloses:

[0037] The present invention contemplates that the determination is made for most source addresses based on an exemplary reference S/D/M/T distribution profile for a non-spoof source address in general. The determination is made using historical S/D/M/T distribution profiles only for a minority number of known non-spoof source addresses, such as known non-spoof source addresses of certain "premium" clients of the network node being "protected".

[0038] Skipping briefly to FIG. 13a-13d and FIG. 14a-14d. FIG. 13a-13b illustrate one each of an example spatial and an example "destination" distribution profile of a source address having spoof instances. Experience has shown that if spoof source addresses are employed in a denial of service attacks against a network node, it is likely that the source addresses will be simultaneously observed in multiple domains of network 100, even domains that are geographically dispersed, as illustrated by the histogram of FIG. 13a. Similarly, if spoof source addresses are employed in a denial of service attacks against a network node, it is likely that the spoof source addresses will not be a subset or substantially related to the source addresses of other packets being routed to other destinations at the routing location, as illustrated by FIG. 13b, where the destinations have disjointed source address ranges for the various destinations of the packets being routed at the routing location. Further, if spoof source addresses are employed in a denial of service attacks against a network node, it is likely that the spoof source addresses will be migrating across different network domains in a very rapid rate, i.e. the routing paths change from one network domain to another relatively quickly, as illustrated by FIG. 13c, having a high number of incidence with short timing duration between routing path changes. Lastly, if spoof source addresses are employed in a denial of service attacks against a network node, it is likely that the source addresses will be repeatedly observed within a very short interval as illustrated by the histogram of FIG. 13b, having an exponentially decay type of profile (in terms of elapsed time between packets with the same source address).

[0039] These characteristics are likely to be different from that of non-spoof source addresses, where spatially, they tend to distribute normally over a domain and its "immediately" adjacent domains, as illustrated by FIG. 14a; and from a destination source address range perspective, they tend to be subset of, or substantially related to source addresses of other packets being routed to other destinations at the routing location, as illustrated by FIG. 14b. From a migration perspective, the number of incidents having short duration between routing path

changes should be very low, as illustrated by FIG. 14c, and from a timing perspective, they too tend to distribute normally over a mean arrival time, as illustrated by FIG. 14d. In addition to being representative of spatial, destination source address range, migration, and timing distribution profiles of a non-spoof source address in general, the S/D/M/T distribution profiles illustrated in FIG. 14a-14d may be actual spatial, destination source address range, migration and timing distribution profiles (historically compiled) of a source address. Such historical profiles may e.g. be compiled for certain premium service clients, as alluded to earlier. Compilation of these exemplary/actual profiles may be performed using any number of statistic gathering techniques known in the art.

Thus, while Wetherall discloses histograms, Wetherall does not suggest much less describe any of the features of ... determining if values of at least one parameter exceed normal, threshold values expected for the parameter to indicate an attack on the site, producing a histogram for the at least one parameter of network traffic to characterize the attack by comparing the histogram to at least one historical histogram for that parameter and filtering out traffic based on characterizing the traffic, which the gateway deems to be part of an attack.

Rather, Wetherall uses the histogram as a template of distribution of source IP addresses to determine that there is a spoofed source IP address attack. Appellant's claimed limitation on the other hand uses the histograms to characterize the type of attack. As discussed above there is no suggestion to combine the references "... to include the filtering processes based on the characterization process using histogram in the invention of Botros in order to route out the undesirable packets as taught in Wetherall see Par. 0041."

Appellant further contends that the motivation advanced by the examiner to combine the references is inadequate for the reasons discussed above.

Claims 33 and 34

The examiner contends that: "Regarding Claim 33-36, Botros discloses the communicating of statistics (sic) to data center over an secure networks see Fig. 3." Appellant disagrees. Claim 33 requires "communicating statistics collected in the gateway to a control center." Botros does not disclose the gateway as discussed above, nor does Botros disclose a control center that receives statistics from the gateway. Claim 34 recites that: "communicating occurs over a dedicated link to the control center via a hardened network." Botros does not specifically disclose any of these features.

Claim 37

Claim 37 is directed to a method to reduce blocking of legitimate traffic in a process to protect a victim site during a denial of service attack. Claim 37 includes the features of producing a histogram of network traffic to characterize an attack and filtering out traffic deemed part of an attack. No combination of Botros with Wetherall discloses the combination of these features as was argued above.

Claim 37 further limits the filter to include the features of constructing a master correlation vector having asserted bits corresponding to the most important parameter correlations, initializing a packet's correlation bit vector to 0, and for every parameter: retrieving the parameter in a parameter suspicious vector to construct the packet's correlation bit vector, and using the value of the packet's correlation bit vector to index into the master correlation bit vector.

The examiner contends that:

23. Regarding Claim 37, Botros discloses the monitoring of device for attacks where detecting process to determine if the parameter of network traffic exceed normal value see Col 10 Ln 40-53 & Col 9 Ln 1-23; a process to build histogram for the parameter see Fig. 9 item 904 & Fig. 10. But does not discloses the filtering out traffic deemed part of attack by producing/retrieving of a vector that is constant/initialized and constructing a vector for packets to test whether to forward for drop packets from source address based on parameter correlations. However, Wetherall discloses the filtering out traffic deemed part of attack by producing of a vector that is constant/initialized to zero and constructing a vector for packets to test/index whether to forward for drop see Par. 0056. It would be obvious to one having ordinary skill in the art at the time of the invention to include the filtering processes based on the characterization process using histogram in the invention of Botros in order to route out the undesirable packets as taught in Wetherall see Par. 0041.

The examiner's argument fails to address any of the specific features of claim 37 involving the aspects of the claimed filtering, namely the features of constructing a master correlation vector and initializing a packet's correlation bit vector and for every parameter, retrieving the parameter in a parameter suspicious vector to construct the packet's correlation bit vector, and using the value of the packet's correlation bit vector to index into the master correlation bit vector.

Claim 38

Claim 38 is allowable over the combination of references since no combination suggests testing the indexed bit in the master correlation vector, where if the bit in the master correlation bit vector is a one, the packet is dropped, otherwise the packet is forwarded.

As discussed above for claim 37, no combination of Botros with Wetherall discloses a master correlation bit vector. Therefore, no combination of Botros with Wetherall could disclose testing of the indexed bit in the master correlation vector and specifically use that technique for dropping or forwarding a packet.

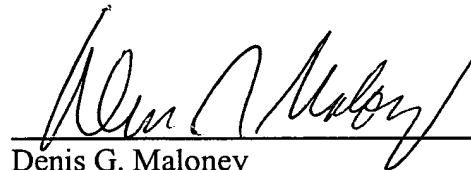
Conclusion

Appellant submits that Claims 7, 9-14, 19-23, 26 are not anticipated under 35 U.S.C. 102(e) by U.S. Patent 6,769, 066 B1 to Botros et al. and that Claims 1-6, 8, 15-18, 24-25, 27, 30-40 are patentable under 35 U.S.C. 103(a) over Botros et al. in view of U.S. Patent Application 2002/0107960 A1 to Wetherall et al. Therefore, the Examiner erred in rejecting Appellant's claims and should be reversed.

Respectfully submitted,

Date: _____

11/2/06



Denis G. Maloney
Reg. No. 29,670

Fish & Richardson P.C.
225 Franklin Street
Boston, MA 02110-2804
Telephone: (617) 542-5070
Facsimile: (617) 542-8906

Appendix of Claims

1. A process that monitors network traffic through a monitoring device disposed between a data center and a network for thwarting denial of service attacks on the data center, the process comprises:

a detection process to determine if the values of a parameter of network traffic exceed normal values for the parameter to indicate an attack on the data center;

a characterization process to build a histogram for the parameter to compute significant outliers in a parameter and classify the attack; and

a filtering process for filtering of network packets based on the characterization process.

2. The process of claim 1 wherein, in the characterization process, suspicious parameter values are represented by a bit vector with a 1 in every position corresponding to a "bad" value, and a 0 in every position corresponding to a "good" value.

3. The process of claim 1 wherein the characterization process comprises:
a correlation process that correlates suspicious parameters and determines existence of correlations of those parameters that indicate types of attacks.

4. The process of claim 3 wherein the correlation process is used to reduce dropping of legitimate traffic.

5. The process of claim 2 wherein filtering is aggregate filtering.

6. The process of claim 1 wherein parameters include at least one of source IP address, destination IP address, source TCP/UDP ports, destination TCP/UDP ports, IP protocol, IP TTL, IP length, hash of payload fragment, IP TOS field, and TCP flags.

7. A method for thwarting denial of service attacks on a data center, the method comprising:

producing a histogram of received network traffic for at least one parameter of network packets; and

characterizing an attack based on comparison of a historical histogram with the produced histogram data for one or more parameters.

8. The method of claim 7 further comprising:

filtering network packets sent to the data center based on whether or not a value of the attribute represented in the current histogram is within a normal range of values for the attribute, as determined by comparison to the historical histogram.

9. The method of claim 7 wherein the historical histogram is based on time periods that can range from 1 hour to 1 week or more.

10. The method of claim 7 wherein the produced histogram is produced during an attack and over time periods of about 10-300 sec or so.

11. The method of claim 7 further comprising:

normalizing the produced and the historical histograms for each parameter; and
computing their difference to identify significant outliers that are considered indicators of suspicious traffic.

12. The method of claim 11 further comprising:

correlating suspicious parameters to reduce blocking of legitimate traffic.

13. The method of claim 12 wherein the bit vector contains sufficient bits to represent the whole parameter space.

14. The method of claim 11 further comprising:
correlating suspicious parameters to determine existence of correlations of those parameters that can point to indications of attacks.
15. The method of claim 11 wherein filtering based on attribute further comprises:
producing a master correlation vector from a stream of sampled packets and examining the network packets using a process that is constant-time, independently of the number of correlations or of the number of suspicious values for a parameter.
16. The method of claim 11 wherein filtering based on attribute further comprises:
constructing a master correlation bit vector corresponding to the most important parameter correlations; and
producing for each packet a correlation bit vector to index into the master correlation bit vector.
17. The method of claim 16 wherein filtering based on attribute further comprises:
testing the bit in the master correlation vector to decide whether to drop or forward the packet.
18. The method of claim 7 wherein attributes include at least one of source IP address, destination IP address, source TCP/UDP ports, destination TCP/UDP ports, IP protocol, IP TTL, IP length, hash of payload fragment, IP TOS field, and TCP flags.
19. The method of claim 7 wherein the method is executed on a data collector.
20. The method of claim 7 wherein the method is executed on a gateway.
21. A monitoring device for thwarting denial of service attacks on a data center, the monitoring device comprises:

a computing device executing:
a process to build at least one histogram for at least one parameter of network traffic; and
a process to characterize an attack based on a comparison of a historical histogram of the
at least one parameter to the built at least one histogram for the at least one parameter.

22. The monitoring device of claim 21 further comprising:
a process to correlate suspicious parameters to reduce blocking of legitimate traffic.

23. The monitoring device of claim 21 wherein the characterization process
normalizes the historical and built histograms for each parameter and computes their difference
to identify significant outliers that are considered indicators of suspicious traffic.

24. The monitoring device of claim 23 wherein the characterization process produces
a master correlation vector from a stream of sampled packets and examines the sampled packets
using a process that is constant-time, independently of the number of correlations or of the
number of suspicious values for a parameter.

25. The monitoring device of claim 21 wherein the device is a gateway device that is
adaptable to dynamically install filters on nearby routers.

26. The monitoring device of claim 21 wherein the device is a data collector.

27. The monitoring device of claim 21 wherein the parameters include at least one of
source IP address, destination IP address, source TCP/UDP ports, destination TCP/UDP ports, IP
protocol, IP TTL, IP length, hash of payload fragment, IP TOS field, and TCP flags.

28. A computer program product residing on a computer readable medium
comprising instructions for causing a processor to:
build a histogram for a parameter of network traffic; and

use the histogram data for the parameter to characterize an attack.

29. The computer program product of claim 28 further comprising instructions to:
filter network traffic based on characterization of the attack.

30. The computer program product of claim 28 further comprising instructions to:
determine if the values of a parameter exceed normal values for the parameter to indicate
an attack on the site;

31. The computer program product of claim 30 further comprising instructions to:
use the histogram to characterize the attack when it is determined that one of the
parameters exceeds a threshold.

32. A method of protecting a data center during a denial of service attack, the method
comprises:

monitoring network traffic through a gateway disposed between the data center and a
network:

determining if values of at least one parameter exceed normal, threshold values expected
for the parameter to indicate an attack on the site;

producing a histogram for the at least one parameter of network traffic to characterize the
attack by comparing the histogram to at least one historical histogram for that parameter; and

filtering out traffic based on characterizing the traffic, which the gateway deems to be
part of an attack.

33. The method of claim 32 further comprising:
communicating statistics collected in the gateway to a control center.

34. The method of claim 33 wherein communicating occurs over a dedicated link to
the control center via a hardened network.

35. The method of claim 33 wherein the gateway is physically deployed in line in the network.

36. The method of claim 33 wherein filtering occurs on nearby routers.

37. A method to reduce blocking of legitimate traffic in a process to protect a victim site during a denial of service attack, comprises:

producing a histogram of network traffic to characterize an attack; and
filtering out traffic deemed part of an attack with filtering comprising:
constructing a master correlation vector having asserted bits corresponding to the most important parameter correlations;
initializing a packet's correlation bit vector to 0, and for every parameter:
retrieving the parameter in a parameter suspicious vector to construct the packet's correlation bit vector; and
using the value of the packet's correlation bit vector to index into the master correlation bit vector.

38. The method of claim 37 further comprising:
testing the indexed bit in the master correlation vector, where if the bit in the master correlation bit vector is a one, the packet is dropped, otherwise the packet is forwarded.

39. The method of claim 37 wherein the master correlation vector is constructed from a stream of sampled packets.

40. The method of claim 37 further comprising:
maintaining a correlation bit vector with as many bits as there are parameters; and
if a parameter's suspicious vector has a 1 in a bit position corresponding to the parameter's value in a packet, the method further comprises:

Applicant : Massimiliano Antonio Poletto et al.
Serial No. : 10/066,232
Filed : January 31, 2002
Page : 32 of 34

Attorney's Docket No.: 12221-010001

setting the bit corresponding to the parameter in the packet's correlation vector to 1.

Applicant : Massimiliano Antonio Poletto et al.
Serial No. : 10/066,232
Filed : January 31, 2002
Page : 33 of 34

Attorney's Docket No.: 12221-010001

Evidence Appendix

None

Applicant : Massimiliano Antonio Poletto et al.
Serial No. : 10/066,232
Filed : January 31, 2002
Page : 34 of 34

Attorney's Docket No.: 12221-010001

Related Proceedings Appendix

None